Consultation response form

Your response

Question	Your response
Question 1: How do you measure the number of users on your service?	Confidential? – Y / N
Question 2: If your service comprises a part on which user-generated content is present and a part on which such content is not present, are you able to distinguish between users of these different parts of the service? If so, how do you make that distinction (including over a given period of time)?	Confidential? – Y / N
Question 3: Do you measure different segments of users on your service? **Do you segment user measurement by different parts of your service? For example, by website vs app, by product, business unit.** **Do you segment user measurement into different types of users? For example: creators, accounts holders, active users.** **How much flexibility does your user measurement system have to define new or custom segments?**	Confidential? – Y / N
Question 4: Do you publish any information about the number of users on your service?	Confidential? – Y / N
Question 5: Do you contribute any user number data to external sources/databases, or help industry measurements systems by tagging or sharing user measurement data? If not, what prevents you from doing so?	Confidential? – Y / N
Question 6: Do you have evidence of functionalities that may affect how easily, quickly and widely content is disseminated on U2U services? **Are there particular functionalities that**	Confidential? – N CST wants to highlight that '*how*

enable content to be disseminated easily on U2U services?

- **Are there particular functionalities that enable content to be disseminated quickly on U2U services?**
- **Are there particular functionalities that enable content to be disseminated widely on U2U services?**
- **Are there particular functionalities that prevent content from being easily, quickly and widely disseminated on U2U services?**

easily, quickly and widely content is disseminated' (often referred to as virality) on U2U services is often facilitated by functionalities that allow users' to easily share content from other users. For example, this may be the 'retweeting' functionality on Twitter/X, the 'share' button on Facebook or the capacity to 'forward' messages on Telegram. These functionalities mean that a single piece of content shared by a user can achieve virality simply based on the number of other users who then share that content. Therefore the more users who then share that content, the more other users on the service are potentially exposed to it – including users who may not wish to see that content because they don't engage with the original source of that content. A CST report into the online drivers of conversations around antisemitism in the Labour Party, found that there were '*36 key pro-Corbyn Twitter accounts, each with their own, overlapping, online networks that drive social media conversations about antisemitism and the Labour Party'*[1]. This relatively small number of accounts was driving the conversation specifically in regard to downplaying or dismissing allegations of antisemitism in the Labour Party. The report also noted that hashtags played a key part in driving online conversations around antisemitism in the Labour Party. For

example:
'According to one analysis of the [Boycott Rachel Riley] *campaign, ten per cent of the accounts using the hashtag were responsible for 45 per cent of the tweets involved.52 This shows the ability of a relatively small number of dedicated Twitter users to generate significant online influence, so that their views and opinions dominate the online conversation and outweigh other trending Twitter content.[2]'*
Hashtags or tagging functionalities therefore may also affect virality and allow a relatively small number of users to easily gain a wider reach, including to non-engaged users who may not wish to view that content.
CST also noted in our 2021 report with the Antisemitism Policy Trust and the Woolf Institute the significant role played by hashtags on Instagram. Specifically, the report noted:
'People exposed to antisemitism on Instagram are not necessarily searching for it. This is a case of antisemitic supply rather than demand. Antisemitic Instagram hashtags, alongside hashtags of antisemitic supply rather than demand. Antisemitic Instagram hashtags, alongside hashtags with demonstrable links to antisemitism, were viewed tens of thousands of times during our with demonstrable links to antisemitism, were viewed tens of thousands of times during our

seven-week research period. Worryingly, the hashtags also generated thousands of likes. seven-week research period[3].'

In this case, hashtags affected, as per the questions set out above, 'how quickly, easily and widely content is disseminated'. The report also noted that hashtags didn't just impact the original source content, but also led users to new types of content. It noted that:

'hashtags that are either antisemitic or commonly that are either antisemitic or commonly associated with antisemitism are often used associated with antisemitism are often used alongside hashtags related to conspiracy alongside hashtags related to conspiracy theories. In the cases we observed, this included theories. In the cases we observed, this included conspiracy theories concerning chemtrails, 5G conspiracy theories concerning chemtrails, 5G and pedophilia[4].'

Therefore, hashtags may also play a role in further dissemination of other types of harmful content, as well as drawing users into digital rabbit holes that may contribute towards other problems such as radicalisation. This can be facilitated by the hashtags themselves, as well as functionalities that, in part, may rely on hashtags to promote content to users. For example, this may be the 'explore' tab on Twitter/X that lists trending

hashtags or the ‘search ‘page on Instagram. In both cases, this impacts how widely content is disseminated on U2U services. The same is true of smaller platforms, such as BitChute who also provide a ‘trending’ page on which they list ‘trending tags’.
In regard to ‘particular functionalities that prevent content from being easily, quickly and widely disseminated on U2U services’, CST notes that frictions in regard to sharing cross-platform links (outlinks) may be an effective tool to prevent virality of content, including illegal and harmful content. This may be especially useful regarding the sharing of links from U2U services that have low or poor content moderation standards. CST notes that Twitter/X used to flag outlinks to BitChute.

Question 7: Do you have evidence relating to the relationship between user numbers, functionalities and how easily, quickly and widely content is disseminated on U2U services?	Confidential? – Y / N
Question 8: Do you have evidence of other objective and measurable factors or characteristics that may be relevant to category 1 threshold conditions?	Confidential? – Y / N
Question 9: Do you have evidence of factors that may affect how content that is illegal or harmful to children is disseminated on U2U services?	Confidential? – Y / N
Are there particular functionalities that play a key role in enabling content that is illegal or harmful to children to be disseminated on U2U services?	

- **Do you have evidence relating to the relationship between user numbers, functionalities and how content that is illegal or harmful to children is disseminated on U2U services?**

Question 10: Do you have evidence of other objective and measurable characteristics that may be relevant to category 2B threshold conditions?

Confidential? – N

CST notes that anonymity continues to play a key role in facilitating harmful content on U2U services. This is an issue that has been raised by CST previously in specific relation to the perpetration of hate crime online, as well as hateful content more generally[5]. CST's research into terrorist content, violent extremism and antisemitism continually sees anonymity as being a key driver behind levels of harm on U2U services. Indeed, a 2020 report by the Antisemitism Policy Trust (APT) specifically highlights that '*those utilising an aggressive discourse choose to remain anonymous online, hiding their true identify for nefarious reasons'*[6]. The role of anonymity online is especially aggravated on U2U services that facilitate a culture of anonymity and on which most users are anonymous. This is the case for services such as 4Chan, noted previously by CST as a service that *'is premised upon complete anonymity: there is no requirement for users to create an account or disclose their identity. Indeed, all users are automatically named as "Anon" on the platform.[7]'.* Ofcom's own report into the Buffalo attack highlights the

role played by 4Chan in radicalising the Buffalo attacker[8]. A report published by the ADL in May 2023 also highlighted the role of anonymity in enabling online harassment and criminality[9]. CST strongly believes that the characteristic of anonymity is a key aggravating factor that should be considered as part of the categorisation process. This needs to be considered alongside other characteristics such as ephemeral messaging that may exacerbate the sharing of harmful content more widely. Other characteristics CST has previously evidenced as increasing certain types of harmful content, relate to user base, business model and governance. This is true in situations when these factors apply in combination. For example, CST has repeatedly raised concerns[10] with regards to the UK video-sharing platform BitChute, who host vast quantities of hateful, antisemitic material. In this case, a potent mix of factors relating to user base, business model and governance contribute to the proliferation of hateful material. The nature of the user base is especially relevant for sites like BitChute that rely on user complaints for identifying and removing hateful content; put simply, if the user base goes to a particular site because they are attracted by hateful content, they are unlikely to report it. This piece[11] from terrorism and extremism expert,

Julia Ebner, highlights how:
'Extremists increasingly bring their followers to alt-tech platforms like Bitchute and Odysee, which then serve as radicalisation rabbit holes, as they host some of the most extreme content while being entirely cut off from outside information.'
The role of these sites to seemingly provide a safe haven for users who have been banned from mainstream platforms increases the risk that those users are utilising those spaces in order to share harmful material. This can include terrorism material, hate and misinformation. This becomes an even more potent risk when paired with poor governance structures. This may be in part due to lack of resource, expertise and/or an unwillingness to impose effective content moderation practices on their service. CST notes that some of these U2U services, like Gab for example, market themselves as '*a social network that champions free speech, individual liberty and the free flow of information online*'. In reality, as noted in several reports,[12] Gab in fact simply serves as a safe haven for hate speech and extremist content. Ofcom must ensure that platforms whose governance model is predicated on a 'free speech' model are not simply using this as cover for hosting hateful and/or extremist content.

Question 11: Do you have evidence of matters that affect the prevalence of content that (once the Bill takes effect) will count as search content that is illegal or harmful to children on particular search services or types of search service? For example, prevalence could refer to the proportion of content surfaced against each search term 16 that is illegal or harmful to children, but we welcome suggestions on additional definitions.

- **Do you have evidence relating to the measurement of the prevalence of content that is illegal or harmful to children on search services?**

Confidential? – N

Search Services have particular functionalities, and these differ depending on whether this is text based or voice-recognition software. In 2021, The Antisemitism Policy Trust, in collaboration with the Community Security Trust (CST), published a study about Google's SafeSearch option and the prevalence of antisemitism in search results. Our research found that its 'SafeSearch' option, which is used by children as it is considered safer, produced as many antisemitic results as its regular search. For example, when searching for the term 'Jew jokes' with the SafeSearch option disabled, 48% of the results produced by Google were found to be antisemitic – a high proportion in of itself. However, the same search phrase with the SafeSearch option enabled, produced an even greater proportion of antisemitic results – 57%. We have

also found that Google's developer tool did not, at the time the research took place, have a specific category for antisemitic, racist or discriminatory content, and therefore could not correctly identify it. Google's public response to the report was that its tool was not designed to filter such content. A remarkable position.

This could place children and other vulnerable people at risk. Parents who allow their children to use Google with the SafeSearch option activated, do so under the assumption that this can limit exposure to harmful content. Unfortunately, our findings show that this is not always the case. While SafeSearch may be effective with regard to certain content, such as pornography, it is less effective when it comes to antisemitic and potentially other racist materials. Since our study was published, Google have remedied some of the results and have addressed antisemitism on their

search platform, although much can still be easily found. However, it shows that vigilance is always needed by search services that are easily accessible to children.

Another concern relevant to search platforms, is the ease of finding social media platforms and websites that host content that is harmful to children. For example, extremist platform BitChute is easily found on regular browsers and a direct link to the website is found on Google even with the ‘enhanced protection’ option turned on. Google defines this option as a ‘proactive protection against dangerous websites…’ However, although BitChute hosts a large volume of harmful and illegal content, including videos or terror attacks watched by far-right extremists, it comes up first in the search.

Other examples of harmful content that can easily be found by children on search engines includes the many pages teaching children how to easily

bypass parental control features on their computer or phone, and how to bypass YouTube's age restrictions. Such content was easily found by us when searching on the commonly-used search engines; Google, Yahoo and Bing.

Voice search technologies present different problems. The Trust highlighted in our responses to calls for evidence by the Independent Commission UK Counter-Terrorism and the DCMS, that voice activated searches also pose a risk. For example, Amazon's Alexa produced an antisemitic conspiracy theory in response to a search. It suggested – based on a single comment posted on Amazon's website – that the Jewish American-Hungarian philanthropist George Soros is responsible for all of the world's evils – a common trope based on antisemitic conspiracies. This is information that has the potential to reach millions of users around the world. Alexa's reach is

extremely wide, and more than 100 million Alexa voice assistants had reportedly been sold worldwide in 2019.

Antisemitic search results come up in other languages too. Last year the Antisemitism Policy Trust found that asking Siri, in Spanish, *"do the Jews control the media?"* prompted a response of articles including details of *"Jewish control international media"* and an article arguing that *"A world famous sociologist claims that the Jews control the media."*

Unlike some of the social media platforms, that restrict (although without much success) the age of its users, search services are easily accessible to all. As such, we recommend that search platforms should produce comprehensive risk assessments, work continually to moderate their search results and address safety issues and harmful information that children may be exposed to.

Question 12: Do you have evidence relating to the number of users on search services and the level of risk of harm to individuals from search content that is illegal or harmful to children?

- **Do you have evidence regarding the relationship between user numbers on search services and the prevalence of search content that is illegal or harmful to children?**

Confidential? – N

In a report about Google published in 2019 by the Antisemitism Policy Trust and the Community Security Trust (CST), we have investigated searches of antisemitic nature in the UK. Searches included violent intentions towards Jews, negative Jewish stereotypes and searches such as '*I hate Jews*' and *'Why are Jews evil?.'*

We found that an average of 170,000 Google searches with antisemitic content is conducted annually in the UK. We also found that search for information on Holocaust denial rise by about 30% on Holocaust Memorial Day each year. Antisemitic conspiracy theories have also enjoyed rising popularity according to our review of Google searches.

Google's auto-complete function was found to have an effect on what content people are exposed to. For

example, when Google removed ‘are Jews evil’ from its auto-complete function in December 2016, searches for this phrase have dropped by 10%. Results generated by search services have a major impact on the content that people are exposed to. It is therefore crucial for search services to proactively direct people away from harmful content and make sure that their auto-complete algorithm does not generate phrases that can lead users towards illegal or harmful materials.

Question 13: Do you have evidence of other objective and measurable characteristics that may be relevant to category 2A threshold conditions?	Confidential? – Y / N

Please complete this form in full and return to os-cfe@ofcom.org.uk.